

This is an Open Access document downloaded from ORCA, Cardiff University's institutional repository:<https://orca.cardiff.ac.uk/id/eprint/132974/>

This is the author's version of a work that was submitted to / accepted for publication.

Citation for final published version:

Ibanez Garcia, Yazmin , Gutierrez Basulto, Victor and Schockaert, Steven 2020. Plausible reasoning about EL-Ontologies using concept interpolation. Presented at: 17th International Conference on Principles of Knowledge Representation and Reasoning (KR 2020), Rhodes, Greece, 12-18 September 2020. Proceedings of the 17th International Conference on Principles of Knowledge Representation and Reasoning. IJCAI Organization, pp. 506-516. 10.24963/kr.2020/51

Publishers page: <https://doi.org/10.24963/kr.2020/51>

Please note:

Changes made as a result of publishing processes such as copy-editing, formatting and page numbers may not be reflected in this version. For the definitive version of this publication, please refer to the published source. You are advised to consult the publisher's version if you wish to cite this paper.

This version is being made available in accordance with publisher policies. See <http://orca.cf.ac.uk/policies.html> for usage policies. Copyright and moral rights for publications made available in ORCA are retained by the copyright holders.



Plausible Reasoning about \mathcal{EL} -Ontologies using Concept Interpolation

Yazmín Ibáñez-García, Víctor Gutiérrez-Basulto, Steven Schockaert

Cardiff University, UK

{ibanezgarciay, gutierrezbasultov, schockaerts1}@cardiff.ac.uk

Abstract

Description logics (DLs) are standard knowledge representation languages for modelling ontologies, i.e. knowledge about concepts and the relations between them. Unfortunately, DL ontologies are difficult to learn from data and time-consuming to encode manually. As a result, ontologies for broad domains are almost inevitably incomplete. In recent years, several data-driven approaches have been proposed for automatically extending such ontologies. One family of methods rely on characterizations of concepts that are derived from text descriptions. While such characterizations do not capture ontological knowledge directly, they encode information about the similarity between different concepts, which can be exploited for filling in the gaps in existing ontologies. To this end, several inductive inference mechanisms have already been proposed, but these have been defined and used in a heuristic fashion. In this paper, we instead propose an inductive inference mechanism which is based on a clear model-theoretic semantics, and can thus be tightly integrated with standard deductive reasoning. We particularly focus on interpolation, a powerful commonsense reasoning mechanism which is closely related to cognitive models of category-based induction. Apart from the formalization of the underlying semantics, as our main technical contribution we provide computational complexity bounds for reasoning in \mathcal{EL} with this interpolation mechanism.

1 Introduction

In the field of AI, knowledge about concepts has traditionally been encoded using logic, often in the form of description logic ontologies (Baader, Horrocks, and Sattler 2004; Baader et al. 2017). While this approach has been highly successful in particular domains, such as health care and biomedical research, the difficulty in acquiring (description logic) ontologies has clearly hampered a more widespread adoption. In open-domain settings, it is almost impossible to exhaustively encode all relevant knowledge about the concepts of interest. As a simple example illustrating this so-called knowledge acquisition bottleneck, the SUMO ontology¹ contains the knowledge that *linguine*, *penne*, *spaghetti*, *couscous* and *ziti* are types of pasta, but none of the many other types of pasta are included.

Beyond the use of ontologies, there has also been a large interest in learning concept representations from data, such

as text descriptions. Word embedding models (Mikolov et al. 2013) learn such representations, for instance. Some authors have also proposed approaches that exploit semi-structured data such as WikiData, Freebase and BabelNet (Neelakantan and Chang 2015; Camacho-Collados, Pilehvar, and Navigli 2016; Jameel, Bouraoui, and Schockaert 2017). Data-driven concept representations are highly complementary to ontologies: they excel at capturing similarity but are otherwise limited in the kinds of dependencies between concepts they can capture. They are essentially tailored towards a form of *inductive reasoning*: given a number of instances of some concept, they are used to predict which other entities are also likely to be instances of that concept. Conversely, traditional ontology languages use rules to encode rigid dependencies between concepts, but they cannot capture graded notions such as similarity, vagueness and typicality. Description logic representations are thus rather tailored to support *deductive reasoning* about concepts.

There is a growing realization that a combination of deductive and inductive reasoning about concepts is needed in many applications (van Harmelen and ten Teije 2019; d’Amato 2020). While several authors have started to explore ways in which such an integration can be achieved, existing work has mostly relied on heuristic methods, focusing on empirical performance rather than the underlying principles. For instance, several approaches have been proposed to exploit rules (Guo et al. 2016; Demeester, Rocktäschel, and Riedel 2016), and symbolic knowledge more generally (Xu et al. 2014; Faruqi et al. 2015), to learn higher-quality vector space representations. Conversely, some authors have used vector representations to infer missing knowledge graph triples (Neelakantan and Chang 2015; Xie et al. 2016), missing ABox assertions (Rizzo et al. 2013; Bouraoui and Schockaert 2018), or missing concept inclusions (Li, Bouraoui, and Schockaert 2019).

The main focus of this paper is on the inference of plausible concepts inclusions, that is, concept inclusions which are not entailed from a given TBox, but which are likely to hold given the knowledge obtained from vector representations and the TBox. However, unlike in previous work, rather than focusing on empirical performance, we aim to study the underlying principles. In particular, in existing approaches, inductive and deductive inferences are typically decoupled. For instance, in (Li, Bouraoui, and Schockaert

¹<http://www.adampease.org/OP/>

2019), missing concept inclusions are predicted in a pre-processing step, after which the standard deductive machinery is employed. The main purpose of this paper is to propose a model-theoretic semantics in which some forms of inductive reasoning about description logic ontologies can be formalized, and which thus allows for a tighter integration between the deductive and inductive inferences.

To illustrate the particular setting that we consider in this paper, assume that we have the following knowledge in a given TBox about some concept C :

$$Rabbit \sqsubseteq C \quad Giraffe \sqsubseteq C$$

If we additionally have background knowledge about rabbits, giraffes and zebras, in particular the fact that zebras satisfy all the *natural* properties that rabbits and giraffes have in common (e.g. being mammals and herbivores), we could then make the following inductive inference, even if we know nothing else about C :

$$Zebra \sqsubseteq C \quad (1)$$

In other words, any natural property that is known to hold for giraffes and rabbits is likely to hold for zebras as well. In such a case, we say that zebras are *conceptually between* rabbits and giraffes. The corresponding inference pattern is known as *interpolation* in AI^2 (Dubois et al. 1997; Schockaert and Prade 2013) and is closely related to the notion of category based induction in cognitive science (Osherson et al. 1990). Importantly, vector representations of concepts, e.g. word embeddings, knowledge graph embeddings, could be used to obtain knowledge about conceptual betweenness. For instance, (Derrac and Schockaert 2015) found that geometric betweenness closely corresponds to conceptual betweenness in vector spaces learned with multi-dimensional scaling.

Apart from conceptual betweenness, the notion of *naturalness* also plays a central role. Indeed, it is clear that the conclusion in (1) can only be justified by making certain assumptions on the concept C . If C could be an arbitrary concept, the resulting inference may clearly not be valid, e.g. this is the case if $C = Rabbit \sqcup Giraffe$ (i.e. *Rabbit* or *Giraffe*). For natural properties, however, interpolative inferences seem intuitively plausible. This idea that only some properties admit inductive inferences has been extensively studied in philosophy, among others by Goodman (1955), who called such properties *projectible*. As an example of a non-projectible property, he introduced the famous example of the property *grue*, which means green up to a given time point and blue afterwards. Along similar lines, Quine (1953) introduced the notion of “natural kinds” to explain why only some properties admit inductive inferences. This notion was developed by Gärdenfors (2000), who introduced the term “natural properties” and suggested that such properties correspond to convex regions in a suitable vector space. To determine which concepts, in a given ontology, are likely to be natural, a useful heuristic is to consider the concept name: concepts that correspond to standard natural language terms

are normally assumed to be natural (Gärdenfors 2014). In this paper, we will simply assume that we are given which concept names are natural.

In particular, we consider the following setting. We are given a standard DL ontology, in addition to a set of conceptual betweenness assertions (i.e. assertions of the form “natural properties that hold for C_1, \dots, C_n should also hold for C ”) and a list of natural concepts. The aim is to reason about the given ontology by combining standard deductive reasoning with the aforementioned interpolation principle. Note that in this way, we maintain a clear separation between deriving knowledge from data-driven representations (i.e. the conceptual betweenness assertions) and the actual reasoning process. We particularly focus on an extension of the description logic \mathcal{EL} (Baader et al. 2017). Our motivation for choosing this logic is its simplicity. \mathcal{EL} allows only for conjunction and existential restrictions as constructors. Nevertheless, \mathcal{EL} has been successfully used for encoding large-scale ontologies like SNOMED CT. Furthermore, reasoning in \mathcal{EL} can be performed in polynomial time. Our proposed extension allows for reasoning with natural concepts and conceptual betweenness, and thus supports interpolative reasoning. Formally defining the semantics of these notions requires an extension to the usual first-order semantics of description logics. Indeed, to capture e.g. that the concept *blue* is natural while *grue* is not, we cannot simply model concepts as sets of individuals. To this end, we consider two alternative approaches for characterizing natural concepts at the semantic level. First, we propose a semantics in which natural concepts are characterized using sets of features. This approach is closely related to formal concept analysis (Wille 1982), and is loosely inspired by the long tradition in cognitive science to model concepts in terms of features (Tversky 1977). Second, we propose a semantics based on vector space representations, inspired by conceptual spaces (Gärdenfors 2000), in which natural concepts correspond to convex regions.

As our main technical contribution, we provide complexity bounds for concept subsumption. Concept subsumption in our considered extension of \mathcal{EL} is CONP-complete under the feature-enriched semantics and PSPACE-hard under the geometric semantics. The difference in complexity between the two proposed semantics intuitively stems from differences in how conceptual betweenness interacts with intersection.

Missing proofs can be found in the extended version of this paper (Ibáñez García, Gutiérrez-Basulto, and Schockaert 2020).

2 Background

We briefly recall some basic notions about description logics, focusing on the \mathcal{EL} logic in particular.

Syntax. Consider countably infinite but disjoint sets of *concept names* N_C and *role names* N_R . These concept and role names are combined to \mathcal{EL} *concepts*, in accordance with the following grammar, where $A \in N_C$ and $r \in N_R$:

$$C, D := \top \mid A \mid C \sqcap D \mid \exists r.C$$

²Not to be confused with the notions of interpolation used to relate logical theories (Craig 1957; Lutz and Wolter 2011)

For instance, $A \sqcap (\exists r.(B \sqcap C))$ is an example of a well-formed \mathcal{EL} concept, assuming $A, B, C \in \mathbf{N}_C$ and $r \in \mathbf{N}_R$. An \mathcal{EL} TBox (ontology) \mathcal{T} is a finite set of concept inclusions (CIs) of the form $C \sqsubseteq D$, where C, D are \mathcal{EL} concepts.

Semantics. The semantics of description logics are usually given in terms of first-order interpretations $(\Delta^{\mathcal{I}}, \cdot^{\mathcal{I}})$. Such interpretations consist of a nonempty domain $\Delta^{\mathcal{I}}$ and an interpretation function $\cdot^{\mathcal{I}}$, which maps each concept name A to a subset $A^{\mathcal{I}} \subseteq \Delta^{\mathcal{I}}$ and each role name r to a binary relation $r^{\mathcal{I}} \subseteq \Delta^{\mathcal{I}} \times \Delta^{\mathcal{I}}$. The interpretation function $\cdot^{\mathcal{I}}$ is extended to complex concepts as follows:

$$\begin{aligned} (\top)^{\mathcal{I}} &= \Delta^{\mathcal{I}}, & (C \sqcap D)^{\mathcal{I}} &= C^{\mathcal{I}} \cap D^{\mathcal{I}}, \\ (\exists r.C)^{\mathcal{I}} &= \{d \in \Delta^{\mathcal{I}} \mid \exists d' \in C^{\mathcal{I}}, (d, d') \in r^{\mathcal{I}}\}. \end{aligned}$$

An interpretation \mathcal{I} satisfies a concept inclusion $C \sqsubseteq D$ if $C^{\mathcal{I}} \subseteq D^{\mathcal{I}}$; it is a model of a TBox \mathcal{T} if it satisfies all CIs in \mathcal{T} . A concept C subsumes a concept D relative to a TBox \mathcal{T} if every model \mathcal{I} of \mathcal{T} satisfies $C \sqsubseteq D$. We denote this by writing $\mathcal{T} \models C \sqsubseteq D$.

3 \mathcal{EL} with In-between and Natural Concepts

We introduce the description logic \mathcal{EL}^{\bowtie} , which extends \mathcal{EL} with *in-between concepts* and *natural concepts*.

Syntax. The main change is that we introduce the in-between constructor, which allows us to describe the set of objects that are between two concepts. Specifically, we write $C \bowtie D$ to denote all objects that are between the concepts C and D . Further, because we will need to differentiate between concepts that are natural and concepts which are not, we will assume that \mathbf{N}_C contains a distinguished infinite set of natural concept names $\mathbf{N}_C^{\text{Nat}}$. The syntax of \mathcal{EL}^{\bowtie} concepts C, D is thus defined by the following grammar, where $A \in \mathbf{N}_C$, $A' \in \mathbf{N}_C^{\text{Nat}}$ and $r \in \mathbf{N}_R$:

$$\begin{aligned} C, D &:= \top \mid A \mid C \sqcap D \mid \exists r.C \mid N \\ N, N' &:= A' \mid N \sqcap N' \mid N \bowtie N' \end{aligned}$$

We will call concepts of the form N, N' *natural concepts*. Notably, we only allow the application of the \bowtie constructor on natural concepts. The reason for this will become clearer once we have defined the semantics. An \mathcal{EL}^{\bowtie} TBox is a finite set of concept inclusions $C \sqsubseteq D$, where C, D are \mathcal{EL}^{\bowtie} concepts.

Example 1. In the following, we will consider the \mathcal{EL}^{\bowtie} TBox \mathcal{T} containing the following concept inclusions:

$$\begin{aligned} \text{Rabbit} &\sqsubseteq \text{Herbivore} & (2) \\ \text{Giraffe} &\sqsubseteq \text{Herbivore} & (3) \\ \text{Zebra} &\sqsubseteq \text{Rabbit} \bowtie \text{Giraffe} & (4) \\ \text{Herbivore} &\sqsubseteq \exists \text{eats.Plant} & (5) \end{aligned}$$

such that $\text{Rabbit}, \text{Zebra}, \text{Giraffe}, \text{Herbivore} \in \mathbf{N}_C^{\text{Nat}}$.

Note that betweenness in the proposed logic \mathcal{EL}^{\bowtie} is modelled using a binary connective. In practice, however, the knowledge we have may relate to more concepts. Indeed, our general aim is to deal with knowledge of the form

“natural properties which hold for all of A_1, \dots, A_k also hold for B ”, or more precisely, that to derive $B \sqsubseteq N$ for a natural concept N , it is sufficient that $A_1 \sqsubseteq N, \dots, A_k \sqsubseteq N$ can be derived. However, in both of the semantics that we consider in this paper, the in-between operator will be associative. For $k \geq 2$, we can thus write $B \sqsubseteq A_1 \bowtie \dots \bowtie A_k$ to encode such knowledge.

Semantics. Our aim is to characterize the semantics of naturalness and betweenness, in accordance with the idea of interpolation. For instance, given a TBox containing the axioms $B \sqsubseteq C \bowtie D$, $C \sqsubseteq N$, $D \sqsubseteq N$ with N a natural concept, we should be able to infer $B \sqsubseteq N$. This means that we need to distinguish between natural concepts and other concepts at the *semantic level*, which is not possible if we simply interpret a concept as a set of objects. We will thus refine the usual first-order interpretations, such that we can characterize (i) which concepts are natural and (ii) which concepts are between which others.

We will consider two possible approaches to define such semantics. First, we will consider *feature-enriched semantics*, which defines a semantics in the spirit of formal concept analysis (Wille 1982). In this case, at the semantic level we associate a set of features with each concept. Note that these features are semantic constructs, which have no direct counterpart at the syntactic level. A concept is then natural if it is completely characterized by these features, while B is between A and C if the set of features associated with B contains the intersection of the sets associated with A and C . Second, we will consider *geometric semantics*, which follows the tradition of Gärdenfors (Gärdenfors 2000). In this case, concepts will be interpreted as regions from a vector space. A concept is then natural if it is interpreted as a convex region, while B is between A and C if the region corresponding with B is geometrically between the regions corresponding with A and C (i.e. in the convex hull of their union). In the following sections, we introduce these two types of semantics in more detail.

4 Feature-Enriched Semantics

In this section we introduce a refinement of the usual first-order interpretations, in which each individual is described using a set of features. Our main motivation here is to find the simplest possible semantics which is rich enough to capture betweenness and naturalness.

4.1 Interpretations

The following definition introduces a refinement of the usual DL interpretations, by introducing features in the spirit of formal concept analysis (FCA).

Definition 1. A feature-enriched interpretation is a tuple $\mathcal{J} = (\mathcal{I}, \mathcal{F}, \pi)$, such that

1. $\mathcal{I} = (\Delta^{\mathcal{I}}, \cdot^{\mathcal{I}})$ is a classical DL interpretation;
2. \mathcal{F} is a non-empty finite set of features;
3. π is a mapping assigning to every element $d \in \Delta^{\mathcal{I}}$ a proper subset of features $\pi(d) \subset \mathcal{F}$;
4. for each proper subset $\mathcal{F}' \subset \mathcal{F}$, there exists an element $d \in \Delta^{\mathcal{I}}$ such that $\mathcal{F}' = \pi(d)$.

The last condition intuitively ensures that the different features are independent, by insisting that every combination of features (apart from \mathcal{F} itself) is witnessed by some individual. The reason why this condition is needed relates to the fact that natural concepts will be characterized in terms of sets of features. For instance, it ensures that two concepts which are characterized by different sets of features cannot have the same extension. Note that only proper subsets of \mathcal{F} are considered, such that we can associate \mathcal{F} itself with the empty concept.

Under \mathfrak{I} , a concept C is interpreted as a pair $C^{\mathfrak{I}} := \langle C^{\mathcal{I}}, \varphi(C) \rangle$ where $C^{\mathcal{I}} \subseteq \Delta^{\mathcal{I}}$ and $\varphi(C)$ is the set of all features from \mathcal{F} which the elements from $C^{\mathcal{I}}$ have in common:

$$\varphi(C) := \bigcap_{d \in C^{\mathcal{I}}} \pi(d).$$

Intuitively, we can think of $\varphi(C)$ as the set of necessary conditions that an individual needs to satisfy to belong to the concept C . The features from \mathcal{F} themselves can thus be seen as a set of primitive conditions that humans might rely on when categorizing individuals. However, note that the considered features do not play any role at the syntactic level, i.e. one cannot directly refer to them and it is not possible to specify them when encoding a TBox.

For standard \mathcal{EL} concepts C , the set $C^{\mathcal{I}}$ is defined as in Section 2. For concepts of the form $N \bowtie N'$, we extend the definition of $\cdot^{\mathcal{I}}$ as follows.

$$(N \bowtie N')^{\mathcal{I}} = \{d \in \Delta^{\mathcal{I}} \mid \varphi(N) \cap \varphi(N') \subseteq \pi(d)\}.$$

Intuitively, $(N \bowtie N')^{\mathcal{I}}$ contains all elements from the domain that have all the features that are common to both N and N' . Finally, each role name r is interpreted as $r^{\mathcal{I}} := r^{\mathcal{I}}$.

Definition 2. Let \mathcal{T} be an \mathcal{EL}^{\bowtie} TBox and $\mathfrak{I} = (\mathcal{I}, \mathcal{F}, \pi)$ a feature-enriched interpretation. We say that \mathfrak{I} is a model of \mathcal{T} , written $\mathfrak{I} \models \mathcal{T}$ if the following hold.

1. $C^{\mathcal{I}} \subseteq D^{\mathcal{I}}$, for every $C \sqsubseteq D$ in \mathcal{T} ;
2. for every natural concept name A in \mathcal{T} , it holds that $A^{\mathcal{I}} = \{d \in \Delta^{\mathcal{I}} \mid \varphi(A) \subseteq \pi(d)\}$.

Example 2. Let us consider the following interpretation $\mathfrak{I} = (\mathcal{I}, \mathcal{F}, \pi)$ of the concept names from Example 1. We define $\mathcal{F} = \{f_1, f_2, f_3, f_4\}$, $\mathcal{I} = (\Delta^{\mathcal{I}}, \cdot^{\mathcal{I}})$ and $\Delta^{\mathcal{I}} = \{d_1, d_2, \dots, d_{15}\}$. Furthermore, π is defined as:

$$\begin{aligned} \pi(d_1) &= \{f_1\} & \pi(d_6) &= \{f_1, f_3\} & \pi(d_{11}) &= \{f_1, f_2, f_3\} \\ \pi(d_2) &= \{f_2\} & \pi(d_7) &= \{f_1, f_4\} & \pi(d_{12}) &= \{f_1, f_2, f_4\} \\ \pi(d_3) &= \{f_3\} & \pi(d_8) &= \{f_2, f_3\} & \pi(d_{13}) &= \{f_1, f_3, f_4\} \\ \pi(d_4) &= \{f_4\} & \pi(d_9) &= \{f_2, f_4\} & \pi(d_{14}) &= \{f_2, f_3, f_4\} \\ \pi(d_5) &= \{f_1, f_2\} & \pi(d_{10}) &= \{f_3, f_4\} & \pi(d_{15}) &= \emptyset \end{aligned}$$

Finally, φ and $\cdot^{\mathcal{I}}$ are defined as follows:

$$\begin{aligned} \varphi(\text{Rabbit}) &= \{f_1, f_2, f_3\} & \text{Rabbit}^{\mathcal{I}} &= \{d_{11}\} \\ \varphi(\text{Zebra}) &= \{f_2, f_3\} & \text{Zebra}^{\mathcal{I}} &= \{d_8, d_{11}, d_{14}\} \\ \varphi(\text{Giraffe}) &= \{f_2, f_3, f_4\} & \text{Giraffe}^{\mathcal{I}} &= \{d_{14}\} \\ \varphi(\text{Herbivore}) &= \{f_3\} & \text{Herbivore}^{\mathcal{I}} &= \{d_3, d_6, d_8, \\ & & & d_{10}, d_{11}, d_{13}, d_{14}\} \end{aligned}$$

It is easy to verify that the conditions from Definition 1 are indeed satisfied by \mathfrak{I} . Furthermore, we have, for instance:

$$\varphi(\text{Rabbit}) \cap \varphi(\text{Giraffe}) = \{f_2, f_3\}$$

and thus

$$(\text{Rabbit} \bowtie \text{Giraffe})^{\mathcal{I}} = \{d_8, d_{11}, d_{14}\}$$

In particular, we have that \mathfrak{I} is a model of (2)–(4). Note that for simplicity, we have not considered (5) in this example (which would require considering more objects and features).

The model we considered in the previous example is rather counter-intuitive, as we would normally think of *Rabbit*, *Zebra* and *Giraffe* as disjoint concepts. However, the feature-enriched semantics is rather restrictive when it comes to modelling disjoint concepts. While this is not an issue for the logic \mathcal{EL} , it shows that a different semantics would be needed for extensions of \mathcal{EL} in which disjointness can be expressed. The geometric semantics, which we discuss in Section 5, is more general in this respect.

4.2 Natural Concepts

Observe that Condition 2 in Definition 2 enforces that the extensions of natural concept names are completely determined by their features. This is indeed in line with the intended semantics of *natural concepts* explained above. This property extends to all natural concepts.

Proposition 1. Let \mathcal{T} be an \mathcal{EL}^{\bowtie} TBox and let $\text{Nat}(\mathcal{T})$ denote the smallest set of concepts such that

- $\top \in \text{Nat}(\mathcal{T})$;
- every concept name $A \in \mathcal{N}_{\mathcal{C}}^{\text{Nat}}$ occurring in \mathcal{T} belongs to $\text{Nat}(\mathcal{T})$;
- if $C, D \in \text{Nat}(\mathcal{T})$, then $C \sqcap D \in \text{Nat}(\mathcal{T})$ and $C \bowtie D \in \text{Nat}(\mathcal{T})$.

Then, for every $C \in \text{Nat}(\mathcal{T})$ and every model $\mathfrak{I} = (\mathcal{I}, \mathcal{F}, \pi)$ of \mathcal{T} it holds that $C^{\mathcal{I}} = \{d \in \Delta^{\mathcal{I}} \mid \varphi(C) \subseteq \pi(d)\}$.

Intuitively, for a natural concept C , its associated set of features $\varphi(C)$ corresponds to necessary and sufficient conditions for an element to belong to the concept. A closely related property of natural concepts is that concept inclusions can be characterized in terms of feature inclusion:

Lemma 1. Let $\mathfrak{I} = (\mathcal{I}, \mathcal{F}, \pi)$ be a feature-enriched interpretation and D a natural concept in \mathfrak{I} . Then for every concept C , $\varphi(D) \subseteq \varphi(C)$ iff $C^{\mathcal{I}} \subseteq D^{\mathcal{I}}$.

4.3 In-between Concepts

Feature-enriched interpretations allow us to define betweenness at the level of the objects in the domain of a given interpretation $(\mathcal{I}, \mathcal{F}, \pi)$. For $d, d_1, d_2 \in \Delta^{\mathcal{I}}$, we will say that d is between d_1 and d_2 , denoted by $\text{bet}(d_1, d, d_2)$, if $\pi(d_1) \cap \pi(d_2) \subseteq \pi(d)$.

Proposition 2. Let $\mathfrak{I} = (\mathcal{I}, \mathcal{F}, \pi)$ be a feature-enriched interpretation. For every pair of natural concepts C, D in \mathfrak{I} such that $C^{\mathcal{I}} \neq \emptyset$ and $D^{\mathcal{I}} \neq \emptyset$, it holds that $(C \bowtie D)^{\mathcal{I}}$ is equal to the following set:

$$\mathbf{B} = \{d \in \Delta^{\mathcal{I}} \mid \exists d_1 \in C^{\mathcal{I}}. \exists d_2 \in D^{\mathcal{I}} \text{ s.t. } \text{bet}(d_1, d, d_2)\}.$$

Observe that **B** provides an intuitive definition of betweenness and that the assumption that C and D are natural is crucial for showing $(C \bowtie D)^{\mathcal{I}} \subseteq \mathbf{B}$. This justifies the syntactic restriction that \bowtie is only applied to natural concepts.

4.4 Link with FCA

There is a clear link between the notion of natural concept in an interpretation $(\mathcal{I}, \mathcal{F}, \pi)$ and the notion of *formal concept* from FCA. Let us consider the formal context $(\Delta^{\mathcal{I}}, \mathcal{F}, \iota)$, where the incidence relation ι is defined as $\iota(d, f)$ iff $f \in \pi(d)$, for $d \in \Delta^{\mathcal{I}}$ and $f \in \mathcal{F}$.

Observation 1. *It holds that C is a natural concept in $(\mathcal{I}, \mathcal{F}, \pi)$ iff $(C^{\mathcal{I}}, \varphi(C))$ is a formal concept of the formal context $(\Delta^{\mathcal{I}}, \mathcal{F}, \iota)$.*

Indeed the following two conditions are satisfied:

$$C^{\mathcal{I}} = \{d \in \Delta^{\mathcal{I}} \mid \iota(d, f) \text{ for all } f \in \varphi(C)\} \quad (6)$$

$$\varphi(C) = \{f \in \mathcal{F} \mid \iota(d, f) \text{ for all } d \in C^{\mathcal{I}}\} \quad (7)$$

Condition (6) follows from the definition of natural concept, while (7) follows from the definition of $\varphi(C)$.

Furthermore, note that Conditions 3 and 4 in Definition 1 ensure that (\emptyset, \mathcal{F}) is also a formal concept, in fact the least element of the concept lattice. In other words, if C is a natural concept in \mathfrak{J} such that $C^{\mathcal{I}} = \emptyset$, then $\varphi(C) = \mathcal{F}$. One consequence of this property is the following.

Observation 2. *For every interpretation $(\mathcal{I}, \mathcal{F}, \pi)$, it holds that $(C \bowtie D)^{\mathcal{I}} \neq \emptyset$ iff $C^{\mathcal{I}} \neq \emptyset$ or $D^{\mathcal{I}} \neq \emptyset$.*

4.5 Interpolation in Feature-enriched Models

The following example illustrates how the feature-enriched semantics enables interpolative inferences.

Example 3. *Consider again the TBox \mathcal{T} from Example 1 and the interpretation \mathfrak{J} from Example 2. It is easy to verify that \mathfrak{J} is a model of \mathcal{T} . Moreover, \mathfrak{J} also satisfies the following concept inclusion:*

$$\text{Zebra} \sqsubseteq \text{Herbivore},$$

which follows the inference pattern explained in the introduction, given that *Herbivore* is a natural concept. In fact, this concept inclusion is entailed by \mathcal{T} . To see this, note that in any model $(\mathcal{J}, \mathcal{F}, \pi)$ of \mathcal{T} , because of the concept inclusions (2) and (3), it holds that

$$\varphi(\text{Herbivore}) \subseteq \varphi(\text{Rabbit}) \quad \varphi(\text{Herbivore}) \subseteq \varphi(\text{Giraffe}).$$

That is,

$$\varphi(\text{Herbivore}) \subseteq \varphi(\text{Rabbit} \bowtie \text{Giraffe}).$$

By Lemma 1, and because of $\text{Herbivore} \in \mathbf{N}_C^{\text{Nat}}$, we have that $(\text{Rabbit} \bowtie \text{Giraffe})^{\mathcal{J}} \subseteq \text{Herbivore}^{\mathcal{J}}$. Finally by concept inclusion (4) we can conclude $\text{Zebra}^{\mathcal{J}} \subseteq \text{Herbivore}^{\mathcal{J}}$.

Clearly, the arguments used in the above example generalize. We thus have the following result, which provides the soundness of *interpolative inferences*.

Lemma 2. *Let \mathcal{T} be an \mathcal{EL}^{\bowtie} TBox, and C, D, B be natural concepts w.r.t \mathcal{T} . If $\mathcal{T} \models \{C \sqsubseteq B, D \sqsubseteq B\}$ then $\mathcal{T} \models C \bowtie D \sqsubseteq B$.*

However, the applicability of this lemma is limited, as it requires specific knowledge about $C \bowtie D$. For example, consider the TBox \mathcal{T}' containing the following assertions:

$$A \sqcap C \sqsubseteq B \quad A \sqcap D \sqsubseteq B \quad X \sqsubseteq C \bowtie D$$

with B a natural concept w.r.t \mathcal{T}' . Since $C \bowtie D$ is characterised by all common features of C and D , a plausible inference from \mathcal{T}' is that $A \sqcap (C \bowtie D) \sqsubseteq B$ holds, which in turns allows us to draw the conclusion that $A \sqcap X \sqsubseteq B$. However, using Lemma 2 we can only soundly infer that $\mathcal{T}' \models (A \sqcap C) \bowtie (A \sqcap D) \sqsubseteq B$, provided that $A \sqcap C$ and $A \sqcap D$ are both natural w.r.t \mathcal{T}' . Thus, we shall investigate under which conditions $A \sqcap (C \bowtie D) \sqsubseteq (A \sqcap C) \bowtie (A \sqcap D)$ holds.

Lemma 3. *Let C, D be natural concepts w.r.t. a given TBox \mathcal{T} . For every model \mathfrak{J} of \mathcal{T} , it holds that $\varphi(C \sqcap D) = \varphi(C) \cup \varphi(D)$.*

We use this property to show that interpolation pattern exemplified above is indeed sound for natural concepts.

Theorem 1. *Let \mathcal{T} be an \mathcal{EL}^{\bowtie} TBox, and let A, B, C, D be natural concepts w.r.t. \mathcal{T} . If $\mathcal{T} \models \{A \sqcap C \sqsubseteq B, A \sqcap D \sqsubseteq B\}$ then $\mathcal{T} \models A \sqcap (C \bowtie D) \sqsubseteq B$.*

5 Geometric Semantics

We now turn to a different approach for defining the semantics of \bowtie and natural concepts, which is inspired by conceptual spaces (Gärdenfors 2000). The main idea is that concepts are represented as regions in a Euclidean space, with natural concepts corresponding to convex regions. One important advantage of the geometric semantics is that it is closer to the vector space embeddings that are commonly used when learning concept representations from data. In other words, if knowledge about conceptual betweenness is learned from vector space representations, then it seems natural to define the semantics in a similar way. Another advantage is that the geometric semantics avoids some of the counter-intuitive restrictions of the feature-enriched semantics, in terms of how betweenness and disjointness interact. This means that the geometric semantics can also be used for extensions of \mathcal{EL} in which disjointness can be expressed, although we leave a detailed study of the computational properties of interpolation in such extensions as a topic for future work. On the other hand, as we will see in the next section, these advantages come at a computational cost, even when staying within the context of \mathcal{EL} .

One key issue of the geometric semantics is that, unlike for the feature-enriched semantics, $X \sqsubseteq C \bowtie D$ does not imply $X \sqcap A \sqsubseteq (C \sqcap A) \bowtie (D \sqcap A)$, even when all of the concepts involved are natural. For this reason, we extend the language of \mathcal{EL}^{\bowtie} with assertions of the form $A \triangleleft (C, D)$, where A is a natural concept and C, D are natural concept names. We will refer to these expressions as *non-interference assertions*. Their aim is to encode how $C \bowtie D$ interacts with intersections with A (explained in

more detailed below). We will refer to the resulting logic as $\mathcal{EL}_{reg}^\infty$. In particular, $\mathcal{EL}_{reg}^\infty$ TBoxes are finite sets of concept inclusions and non-interference assertions.

5.1 Interpretations

Geometric interpretations represent concepts as regions, where individuals are intuitively represented as points. In addition to specifying these regions, however, geometric interpretations also specify some additional mappings, which will be needed to formalize the idea of non-interference.

In what follows, we use $\text{conv}(X)$ to denote the *convex hull* of X , that is the intersection of all the convex sets that contain X , and \oplus to denote the concatenation of vectors.

Definition 3 (Geometric interpretation). *Let $\Sigma \subset \mathbf{N}_C \cup \mathbf{N}_R$. An m -dimensional geometric Σ -interpretation \mathcal{I} assigns to every concept name $A \in \Sigma$ a region $\text{reg}_{\mathcal{I}}(A) \subseteq \mathbb{R}^m$ and to every role $r \in \Sigma$ a region $\text{reg}_{\mathcal{I}}(r) \subseteq \mathbb{R}^{2 \cdot m}$. Furthermore, \mathcal{I} specifies for all natural concept names A, B in $\Sigma \cap \mathbf{N}_C^{\text{Nat}}$, a mapping $\kappa_{(A,B)}^{\mathcal{I}}$ from $\text{conv}(\text{reg}_{\mathcal{I}}(A) \cup \text{reg}_{\mathcal{I}}(B))$ to $\text{reg}_{\mathcal{I}}(A) \times \text{reg}_{\mathcal{I}}(B)$ such that for every $p \in \text{conv}(\text{reg}_{\mathcal{I}}(A) \cup \text{reg}_{\mathcal{I}}(B))$ with $\kappa_{(A,B)}^{\mathcal{I}}(p) = (p_1, p_2)$ it holds that*

- p is between p_1 and p_2 , i.e. $p = \lambda p_1 + (1 - \lambda)p_2$ for some $\lambda \in [0, 1]$;
- $\kappa_{(B,A)}^{\mathcal{I}}(p) = (p_2, p_1)$.

Note that for a point $p \in \text{conv}(\text{reg}_{\mathcal{I}}(A) \cup \text{reg}_{\mathcal{I}}(B))$ it is always possible to find points $p_1 \in \text{reg}_{\mathcal{I}}(A)$ and $p_2 \in \text{reg}_{\mathcal{I}}(B)$ such that p is between p_1 and p_2 . Intuitively, however, the mapping $\kappa_{(A,B)}^{\mathcal{I}}$ selects the pair (p_1, p_2) which is most “similar” to p . This intuition will be made explicit when discussing the semantics of non-interference assertions below.

The interpretation of complex \mathcal{EL}^∞ concepts is defined as follows.

$$\begin{aligned} \text{reg}_{\mathcal{I}}(\top) &= \mathbb{R} \\ \text{reg}_{\mathcal{I}}(C \sqcap D) &= \text{reg}_{\mathcal{I}}(C) \cap \text{reg}_{\mathcal{I}}(D) \\ \text{reg}_{\mathcal{I}}(\exists r.C) &= \{p \in \mathbb{R}^m \mid \exists p' \in \text{reg}_{\mathcal{I}}(C), p \oplus p' \in \text{reg}_{\mathcal{I}}(r)\} \\ \text{reg}_{\mathcal{I}}(C_1 \bowtie C_2) &= \text{conv}(\text{reg}_{\mathcal{I}}(C_1) \cup \text{reg}_{\mathcal{I}}(C_2)) \end{aligned}$$

Note how the definition of $\text{reg}_{\mathcal{I}}(C_1 \bowtie C_2)$ defines conceptual betweenness in terms of geometric betweenness, i.e. the instances of the concept $C_1 \bowtie C_2$ are intuitively those individuals which are modelled by points that are geometrically between the regions modelling C_1 and C_2 .

Example 4. Figure 1 depicts a two-dimensional Σ -interpretation of the concepts Rabbit, Zebra, Giraffe and Herbivore from Example 1.

The semantics of $\mathcal{EL}_{reg}^\infty$ TBox assertions is defined as follows. An m -dimensional Σ -interpretation \mathcal{I} satisfies a concept inclusion $C \sqsubseteq D$, for $C, D \in \Sigma$, if $\text{reg}_{\mathcal{I}}(C) \subseteq \text{reg}_{\mathcal{I}}(D)$. The interpretation \mathcal{I} satisfies the non-interference assertion $X \triangleleft (A, B)$ if for all $p \in (X \sqcap (A \bowtie B))^{\mathcal{I}}$, whenever $\kappa_{(A,B)}^{\mathcal{I}}(p) = (p_1, p_2)$, it holds that $p_1 \in X^{\mathcal{I}}$.

The intuition behind non-interference relies on the notion of domains from the theory of conceptual spaces. For instance, if $X \sqsubseteq \text{Red} \bowtie \text{Blue}$ then we would expect that

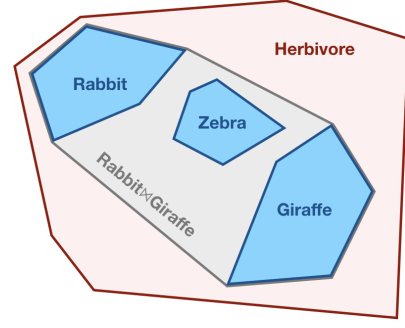


Figure 1: Illustration of a two-dimensional geometric Σ -interpretation.

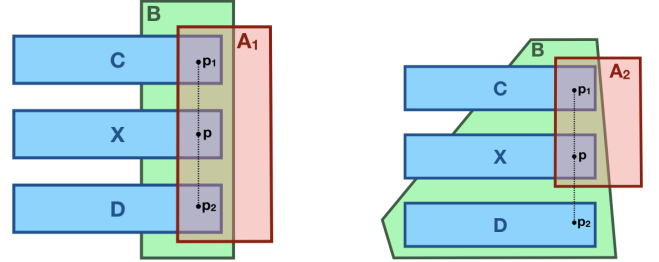


Figure 2: A configuration in which both $A_1 \triangleleft (C, D)$ and $A_1 \triangleleft (D, C)$ can be satisfied (left), and a configuration in which $A_2 \triangleleft (C, D)$ can be satisfied but not $A_2 \triangleleft (D, C)$ (right)

$(X \sqcap \text{Small}) \sqsubseteq (\text{Red} \sqcap \text{Small}) \bowtie (\text{Blue} \sqcap \text{Small})$ also holds. This is because *Red* and *Blue* are defined in the color domain, whereas *Small* is defined in the size domain. Concepts that rely on disjoint sets of domains intuitively cannot interfere with betweenness assertions. In our setting, we only have a single vector space, whereas in the theory of conceptual spaces each domain corresponds to a separate vector space. Instead, as was argued in (Jameel and Schockaert 2016), we can think of such domains as sub-spaces of \mathbb{R}^m . The intended intuition is that $\kappa_{(A,B)}^{\mathcal{I}}(p)$ selects a pair of points (p_1, p_2) which only differ from p in the sub-spaces of the domains that are relevant to A and B . The statement $X \triangleleft (A, B)$ then intuitively asserts that the domains that are relevant for A and B are disjoint from the domains that are relevant for X .

Example 5. The left-hand side of Figure 2 depicts a configuration in which the non-interference assertions $A_1 \triangleleft (C, D)$ and $A_1 \triangleleft (D, C)$ can be satisfied. In particular, a suitable mapping $\kappa_{(A,B)}^{\mathcal{I}}(p) = (p_1, p_2)$ can be found by choosing $p_1 \in \text{reg}_{\mathcal{I}}(C)$ and $p_2 \in \text{reg}_{\mathcal{I}}(D)$ such that p_1 and p_2 share their first coordinate with p . Similarly, the right-hand side of Figure 2 shows a configuration in which $A_2 \triangleleft (C, D)$ can be satisfied, but not $A_2 \triangleleft (D, C)$.

Definition 4. Let \mathcal{T} be an $\mathcal{EL}_{reg}^\infty$ TBox. An m -dimensional geometric Σ -interpretation \mathcal{I} is an m -dimensional Σ -model of \mathcal{T} if the following are satisfied:

1. all the concept and role names appearing in \mathcal{T} are included in Σ ;

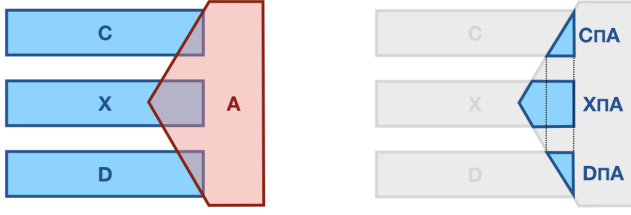


Figure 3: Even though the region representing X is between those representing C and D (left), the region for $X \sqcap A$ is not between those for $C \sqcap A$ and $D \sqcap A$ (right).

2. \mathcal{I} satisfies every concept inclusion in \mathcal{T} ;
3. \mathcal{I} satisfies every non-interference assertion in \mathcal{T} ;
4. $\text{reg}_{\mathcal{I}}(A)$ is a convex region for every concept name $A \in \Sigma \cap \mathcal{N}_{\mathcal{C}}^{\text{Nat}}$.

We refer to \mathcal{I} as a *geometric model*, or simply a *model*, if m and Σ are clear from the context.

5.2 Interpolation in Geometric Models

It is easy to see that Lemma 2 also holds for the geometric semantics, showing that basic interpolative inferences are sound. The reason why we need to consider non-interference is related to the interaction between \bowtie and \sqcap . Recall from Section 4 that interpolation with inclusions of the form $A \sqcap C \sqsubseteq B$ and $A \sqcap D \sqsubseteq B$ required A, C and D to be natural concepts in order to have e.g., that $\varphi(A \sqcap C) = \varphi(A) \cup \varphi(C)$. In the case of the geometric semantics, naturalness alone is not sufficient to allow us to derive $\mathcal{T} \models A \sqcap (C \bowtie D) \sqsubseteq B$ from $\mathcal{T} \models \{A \sqcap C \sqsubseteq B, A \sqcap D \sqsubseteq B\}$. To see this, consider the 2-dimensional interpretation illustrated in Figure 3. We have that $\text{reg}_{\mathcal{I}}(X)$ is between $\text{reg}_{\mathcal{I}}(C)$ and $\text{reg}_{\mathcal{I}}(D)$ and that $\text{reg}_{\mathcal{I}}(C \sqcap A)$ and $\text{reg}_{\mathcal{I}}(D \sqcap A)$ are convex (which is the geometric characterization of naturalness). Nonetheless, we can see that $\text{reg}_{\mathcal{I}}(X \sqcap A)$ is not between $\text{reg}_{\mathcal{I}}(C \sqcap A)$ and $\text{reg}_{\mathcal{I}}(D \sqcap A)$. A way to enable interpolative reasoning between conjunctions of concepts is to require that the concepts involved are non-interfering. This is formalized in the following lemma.

Proposition 3. *Let \mathcal{T} be an $\mathcal{EL}_{\text{reg}}^{\bowtie}$ -TBox. If $\mathcal{T} \models \{A \sqsubseteq (C, D), A \sqcap C \sqsubseteq B, D \sqsubseteq B\}$, with A, B natural concepts and $C, D \in \mathcal{N}_{\mathcal{C}}^{\text{Nat}}$, then $\mathcal{T} \models A \sqcap (C \bowtie D) \sqsubseteq B$.*

Analogously, we can show the following result.

Proposition 4. *Let \mathcal{T} be an $\mathcal{EL}_{\text{reg}}^{\bowtie}$ -TBox. If $\mathcal{T} \models \{A \sqsubseteq (C, D), A \sqsubseteq (D, C), A \sqcap C \sqsubseteq B, A \sqcap D \sqsubseteq B\}$, with A, B natural concepts and $C, D \in \mathcal{N}_{\mathcal{C}}^{\text{Nat}}$, then $\mathcal{T} \models A \sqcap (C \bowtie D) \sqsubseteq B$.*

Example 6. *The configurations in Figure 2 illustrate how sound interpolative inferences can be made when the conditions from Proposition 3 (right-hand side) or Proposition 4 (left-hand side) are satisfied.*

Finally, the following result shows that non-interference is closed under intersection.

Proposition 5. *Let A, B be natural concepts and C, D natural concept names. If $\mathcal{T} \models A \sqsubseteq (C, D)$ and $\mathcal{T} \models B \sqsubseteq (C, D)$ then $\mathcal{T} \models (A \sqcap B) \sqsubseteq (C, D)$.*

6 Complexity of Reasoning with Interpolation

We next analyze the computational complexity of reasoning in \mathcal{EL}^{\bowtie} and $\mathcal{EL}_{\text{reg}}^{\bowtie}$. We show that the ability to perform interpolation increases the complexity of concept subsumption relative to a TBox.

6.1 Concept Subsumption in \mathcal{EL}^{\bowtie}

We start by studying \mathcal{EL}^{\bowtie} and establish that concept subsumption relative to \mathcal{EL}^{\bowtie} -TBoxes is CONP-complete. We show hardness by reducing non-entailment in propositional logic to concept subsumption in \mathcal{EL}^{\bowtie} . The main underlying idea is that a concept inclusion of the form:

$$X_1 \sqcap \dots \sqcap X_n \sqsubseteq Y_1 \bowtie \dots \bowtie Y_m$$

can be used to simulate a propositional clause of the following form:

$$\neg y_1 \vee \dots \vee \neg y_m \vee x_1 \vee \dots \vee x_n$$

where each atom x_i or y_i is associated with a natural concept name X_i or Y_i . This correspondence allows us to reduce the problem of entailment checking in propositional logic to the problem of checking concept subsumption relative to \mathcal{EL}^{\bowtie} -TBoxes; the proof can be found in the online appendix.

Theorem 2. *Concept subsumption relative to an \mathcal{EL}^{\bowtie} -TBox is CONP-hard, even when restricting to TBoxes without any occurrences of existential restrictions.*

For the matching upper bound we provide a polynomial time guess-and-check procedure. We assume that \mathcal{EL}^{\bowtie} -TBoxes are in the following normal form. For a TBox to be in normal form, we require that every concept inclusion is of one of the forms $A \sqsubseteq B$, $A_1 \sqcap A_2 \sqsubseteq B$, $A \sqsubseteq \exists r.B$, $\exists r.A \sqsubseteq B$, $A \sqsubseteq B_1 \bowtie B_2$, $B_1 \bowtie B_2 \sqsubseteq A$, where A, A_1, A_2, B are concept names or the concept \top and B_1, B_2 are natural concept names. It is standard to show that every TBox can be transformed into this normal form in polynomial time such that (non-)subsumption between the concept names that occur in the original TBox is preserved.

We start by showing the following property of feature-enriched interpretations.

Lemma 4. *Let \mathcal{T} be an \mathcal{EL}^{\bowtie} TBox. For every model $\mathcal{I} = (\mathcal{I}, \mathcal{F}, \pi)$ of \mathcal{T} , there is a model $\hat{\mathcal{I}} = (\mathcal{I}, \hat{\mathcal{F}}, \hat{\pi})$ such that $|\hat{\mathcal{F}}| \leq \text{poly}(\mathcal{T})$.*

Before presenting the decision procedure, we introduce some notions. Let \mathcal{T} be an \mathcal{EL}^{\bowtie} TBox. A *feature assignment* for \mathcal{T} from a set of features \mathcal{F} is a mapping θ assigning to each concept name in \mathcal{T} a subset $F \subseteq \mathcal{F}$. We say that a feature assignment θ for \mathcal{T} is *proper* if the following conditions hold:

1. For every concept inclusion of the form $A_1 \sqcap A_2 \sqsubseteq B$ in \mathcal{T} with A_1, A_2 natural concept names, it holds that $\hat{\theta}(B) \subseteq \theta(A_1) \cup \theta(A_2)$;
2. for every concept inclusion of the form $A \sqsubseteq C$ in \mathcal{T} , it holds that $\hat{\theta}(C) \subseteq \theta(A)$,

with $\widehat{\theta}(\cdot)$ defined as follows: $\widehat{\theta}(A) = \theta(A)$;

$$\widehat{\theta}(A \bowtie B) = \theta(A) \cap \theta(B); \quad \widehat{\theta}(\top) = \widehat{\theta}(\exists r.B) = \emptyset.$$

We are now ready to describe our guess-and-check procedure to decide non-subsumption in \mathcal{EL}^{\bowtie} . Given an \mathcal{EL}^{\bowtie} TBox \mathcal{T} , we proceed as follows:

1. Guess a feature assignment θ for \mathcal{T} from some set of features \mathcal{F} (By Lemma 4, we can assume that $|\mathcal{F}| \leq \text{poly}(|\mathcal{T}|)$).
2. Add concept inclusions $A \sqsubseteq C$ to \mathcal{T} if $\widehat{\theta}(C) \subseteq \theta(A)$, for $A \in \mathcal{N}_C$ and $C \in \mathcal{N}_C^{\text{Nat}}$ or $C = B_1 \bowtie B_2$, occurring in \mathcal{T} . Let \mathcal{T}_θ be the TBox obtained after this step.
3. Compute the completion \mathcal{T}'_θ of \mathcal{T}_θ using the classical completion algorithm for \mathcal{EL} (Baader et al. 2017), where concepts of the form $A \bowtie B$ are regarded as concept names. Let \mathcal{T}'_θ be the TBox obtained after this step.
4. Check that θ is proper for \mathcal{T}'_θ .

Lemma 5. *Let \mathcal{T} be an \mathcal{EL}^{\bowtie} TBox and A, B concept names. Then, $\mathcal{T} \not\models A \sqsubseteq B$ iff, after applying Steps 1-4 above, $A \sqsubseteq B \notin \mathcal{T}'_\theta$.*

Summing up we obtain the following result.

Theorem 3. *Concept subsumption relative to \mathcal{EL}^{\bowtie} -TBoxes is CONP-complete.*

6.2 Concept Subsumption in $\mathcal{EL}^{\bowtie}_{\text{reg}}$

We move now to investigate concept subsumption relative to $\mathcal{EL}^{\bowtie}_{\text{reg}}$ TBoxes (under the geometric semantics). In this case, we are able to show the following hardness result.

Theorem 4. *Concept subsumption relative to $\mathcal{EL}^{\bowtie}_{\text{reg}}$ -TBoxes is PSPACE-hard.*

Inspired by (Schockaert and Prade 2013), the proof proceeds by reduction from the dominance problem in generalized CP-nets (GCP-nets). We briefly sketch this reduction. First we recall the dominance problem. A GCP-net over a set of propositional atoms $\mathcal{A} = \{a_1, \dots, a_m\}$ is specified by a set of so-called *conditional preference (CP) rules* ρ_i ($i \in \{1, \dots, n\}$) of the following form:

$$\eta_i : q_i > \neg q_i \quad (8)$$

where η_i is a conjunction of literals and q_i is a literal (over \mathcal{A}). The intuition of this rule is that whenever η_i is true, then it is better to have q_i true than to have q_i false, assuming that everything else stays the same (i.e. *ceteris paribus*). An *outcome* is defined as a tuple of literals (l_1, \dots, l_m) where l_i is either a_i or $\neg a_i$. Outcomes thus encode possible worlds. Let $\omega_1 = (l_1, \dots, \neg q_i, \dots, l_m)$ and $\omega_2 = (l_1, \dots, q_i, \dots, l_m)$ be outcomes which only differ in the truth value they assign to q_i and which both satisfy the condition η_i . Then we say that the rule ρ_i *sanctions an improving flip* from ω_1 to ω_2 . Moreover, we say that an outcome ω *dominates* an outcome ω' , written $\omega' \prec \omega$ if there exists a sequence of improving flips from ω' to ω . A GCP-net is *consistent* if there are no cycles of improving flips, i.e. there are no outcomes ω for which $\omega \prec \omega$. It was shown in (Goldsmith et al. 2008) that the problem of checking whether some outcome ω dominates

an outcome ω' is PSPACE-complete, even when restricted to consistent GCP-nets.

The proposed reduction is defined as follows. Let the initial outcome be given by (l_1, \dots, l_m) . The corresponding $\mathcal{EL}^{\bowtie}_{\text{reg}}$ TBox \mathcal{T} contains the following corresponding concept inclusion:

$$\tau(l_1) \sqcap \dots \sqcap \tau(l_m) \sqsubseteq Z \quad (9)$$

where the mapping τ is defined by $\tau(a_i) = A_i$ and $\tau(\neg a_i) = \overline{A_i}$, with $A_i, \overline{A_i}$ and Z natural concept names. We furthermore extend this mapping to conjunctions of literals as $\tau(l_1 \wedge \dots \wedge l_k) = \tau(l_1) \sqcap \dots \sqcap \tau(l_k)$. For each CP-rule ρ_i , we have that \mathcal{T} contains the following concept inclusions:

$$X_i \sqsubseteq Z \quad (10)$$

$$\tau(\eta_i \wedge q_i) \sqsubseteq W_i \bowtie X_i \quad (11)$$

$$W_i \sqsubseteq \tau(\eta_i \wedge \neg q_i) \quad (12)$$

$$\tau(\eta_i \wedge \neg q_i) \sqsubseteq W_i \quad (13)$$

where W_i and X_i are natural concept names. Furthermore, for each rule $\eta_i : q_i > \neg q_i$ and each a_j , such that neither a_j nor $\neg a_j$ occurs in η_i or $\{q_i, \neg q_i\}$, we add the non-interference assertions $A_j \prec (W_i, X_i)$ and $\overline{A_j} \prec (W_i, X_i)$.

Then we can show that (r_1, \dots, r_m) dominates (l_1, \dots, l_m) , with $(r_1, \dots, r_m) \neq (l_1, \dots, l_m)$ iff

$$\mathcal{T} \models \tau(r_1 \wedge \dots \wedge r_m) \sqsubseteq Z$$

We conjecture that a matching PSPACE upper bound can be found, although this currently remains an open question.

7 Related Work

The problem of automated knowledge base completion has received significant attention in recent years. Most of the work in this area has focused on completing knowledge graphs by learning a suitable vector space representation of the entities and relations involved (Bordes et al. 2013). However, some work has also focused on description logic ontologies. An early example is (d'Amato et al. 2009), which proposed to use a similarity metric between individuals to find plausible answers to queries. More recently, Bouraoui and Schockaert (2018) proposed a method for finding plausible missing ABox assertions, by representing each concept as a Gaussian distribution in a vector space, while Kulmanov et al. (2019) proposed a method to learn a vector space embedding of \mathcal{EL} ontologies for this purpose. The problem of completing TBoxes using vector space representations was considered in (Bouraoui and Schockaert 2019).

The previously mentioned approaches are essentially heuristic, focusing on the empirical performance of the considered strategies, without introducing a corresponding model-theoretic semantics or studying the formal properties of associated reasoning tasks (e.g. computational complexity). The problem of formally combining logics and similarity is addressed in (Sheremet et al. 2005; Sheremet, Wolter, and Zakharyashev 2010), where an operator is introduced to express that a concept A is more similar to some concept B than to some concept C . By focusing

on comparative similarity, the problem of dealing with numerical degrees is avoided. We can thus think of comparative similarity and conceptual betweenness as two complementary approaches for reasoning about similarity in a qualitative way. Related to concept betweenness is the notion of *least common subsumer* (LCS) which has been broadly studied in the context of DLs as means for supporting inductive inference (Cohen, Borgida, and Hirsh 1992; Küsters and Borgida 2001; Baader, Sertkaya, and Turhan 2007; Ecke and Turhan 2012; Zarri   and Turhan 2013; Jung, Lutz, and Wolter 2020). Similarly to the LCS of A and B , $lcs(A, B)$, $A \bowtie B$ subsumes both A and B , thus generalizing them. However, $lcs(A, B)$ is minimal w.r.t. the extensions of A and B , whereas $A \bowtie B$ is minimal w.r.t. their intent under the feature-based semantics, i.e., it is the least common ‘natural subsumer’. The latter is arguably, closer to the cognitive notion of the least common subsumer of A and B as the concept capturing their commonalities. Further, in \mathcal{EL} , where a syntactic description of LCS is not guaranteed to exist, betweenness provides such description. Beyond qualitative approaches, it is also possible to directly model degrees of similarity. For instance, Esteva et al. (1997) considered a graded modal logic which formalizes a form of similarity based reasoning. Fuzzy and rough description logics can also be viewed from this angle (Straccia 2001; Bobillo et al. 2015; Schlobach, Klein, and Peelen 2007; Klein, Mika, and Schlobach 2007; Jiang et al. 2009; Pe  aloza and Zou 2013; Lisi and Straccia 2013). Within a broader context, (Lieto and Pozzato 2018) is also motivated by the idea of combining description logics with ideas from cognitive science, although their focus is on modelling typicality effects and compositionality, e.g. inferring the meaning of *pet fish* from the meanings of *pet* and *fish*, which is a well-known challenge for cognitive systems since typical pet fish are neither typical pets nor typical fish.

The idea of providing a semantics for description logics in which concepts correspond to convex regions from a conceptual space was already considered in (  z  ep and M  ller 2013). Guti  rrez-Basulto and Schockaert (2018) also studied a semantics based on conceptual spaces for existential rules. The idea of linking description logic concepts to feature based models has been previously considered as well. For instance, Porello et al. (2019) introduced a syntactic construct to define description logic concepts in terms of weighted combinations of properties, although unlike in this paper, their properties/features are also syntactic objects. Description logics with *concrete domains* provide means to refer to concrete objects, such as numbers or spatial regions (Lutz 2002; Lutz and Milicic 2007), but (unlike in our case) they come equipped with syntax to access and impose constraints on these domains.

The study of the link between DLs and formal concept analysis has received considerable attention, see e.g. (Baader and Molitor 2000; Rudolph 2006; Baader et al. 2007; Sertkaya 2010; Distel 2011) and references therein. However, unlike in these works, our main objective in this paper is to use features to characterize natural concepts, and provide semantics capturing the idea of interpolation.

8 Conclusions and Future Work

Our central aim in this paper was to formalize interpolative reasoning, a commonsense inference pattern that underpins a cognitively plausible model of induction, in the context of description logics. To this end, we have studied extensions of the description logic \mathcal{EL} in which we can encode that one concept is between two other concepts. In particular, we have studied two approaches to formally define the semantics of betweenness and the related notion of naturalness: one inspired by formal concept analysis and one inspired by conceptual spaces. We furthermore showed that reasoning in the considered extensions of \mathcal{EL} is CONP-complete under the featured-enriched semantics, and PSPACE-hard under the geometric semantics.

There are several important avenues for future work. First, at the foundational level, we believe that our framework can be used as a basis for integrating inductive and deductive reasoning more broadly. Essentially, inductive reasoning requires two things: (i) we need knowledge about the representation of concepts in a suitable feature space and (ii) we need to make particular restrictions on how concepts are represented in that feature space. In this paper, (i) was addressed by providing knowledge about conceptual betweenness whereas (ii) was addressed by the notion of naturalness. However, there may be several other mechanisms to encode knowledge about the feature space. One possibility is to have assertions that relate to analogical proportions (i.e. assertions of the form “ a is to b what c is to d ”), which can be formalized in terms of discrete features or geometric representations (Miclet, Bayoudh, and Delhay 2008; Prade and Richard 2014).

Another important line for future work relates to applying the proposed framework in practice, e.g. for ontology completion or for plausible query answering. This will require two additional contributions. First, we either need a tractable fragment of the considered logics or an efficient approximate inference technique. Second, we need practical mechanisms to deal with the noisy nature of the available knowledge about betweenness (which typically would be learned from data) and the inconsistencies that may arise from applying interpolation (e.g. because concepts that were assumed to be natural may not be). To this end, we plan to study probabilistic or non-monotonic extensions of our framework.

Acknowledgments

Yazm  n Ib    ez-Garc  a and Steven Schockaert have been supported by ERC Starting Grant 637277.

References

- Baader, F., and Molitor, R. 2000. Building and structuring description logic knowledge bases using least common subsumers and concept analysis. In *ICCS*, volume 1867 of *Lecture Notes in Computer Science*, 292–305. Springer.
- Baader, F.; Ganter, B.; Sertkaya, B.; and Sattler, U. 2007. Completing description logic knowledge bases using formal concept analysis. In *IJCAI*, 230–235.

- Baader, F.; Horrocks, I.; Lutz, C.; and Sattler, U. 2017. *An Introduction to Description Logic*. Cambridge Uni. Press.
- Baader, F.; Horrocks, I.; and Sattler, U. 2004. Description logics. In *Handbook on Ontologies*. Springer. 3–28.
- Baader, F.; Sertkaya, B.; and Turhan, A. 2007. Computing the least common subsumer w.r.t. a background terminology. *J. Applied Logic* 5(3):392–420.
- Bobillo, F.; Cerami, M.; Esteva, F.; García-Cerdaña, À.; Peñaloza, R.; and Straccia, U. 2015. Fuzzy description logic. In *Handbook of Mathematical Fuzzy Logic Volume 3*, volume 58 of *Studies in Logic*. College Publications.
- Bordes, A.; Usunier, N.; Garcia-Duran, A.; Weston, J.; and Yakhnenko, O. 2013. Translating embeddings for modeling multi-relational data. In *Proc. of NIPS*. Curran Associates, Inc. 2787–2795.
- Bouraoui, Z., and Schockaert, S. 2018. Learning conceptual space representations of interrelated concepts. In *Proc. of IJCAI*, 1760–1766.
- Bouraoui, Z., and Schockaert, S. 2019. Automated rule base completion as bayesian concept induction. In *Proc. of AAAI*.
- Camacho-Collados, J.; Pilehvar, M. T.; and Navigli, R. 2016. Nasari: Integrating explicit knowledge and corpus statistics for a multilingual representation of concepts and entities. *Artificial Intelligence* 240:36–64.
- Cohen, W. W.; Borgida, A.; and Hirsh, H. 1992. Computing least common subsumers in description logics. In *AAAI*, 754–760. AAAI Press / The MIT Press.
- Craig, W. 1957. Three uses of the herbrand-gentzen theorem in relating model theory and proof theory. *J. Symb. Log.* 22(3):269–285.
- d’Amato, C.; Fanizzi, N.; Fazzinga, B.; Gottlob, G.; and Lukasiewicz, T. 2009. Combining semantic web search with the power of inductive reasoning. In *Proc. of URSW*, 15–26.
- d’Amato, C. 2020. Machine learning for the semantic web: Lessons learnt and next research directions. *Semantic Web* 11(1):195–203.
- Demeester, T.; Rocktäschel, T.; and Riedel, S. 2016. Lifted rule injection for relation embeddings. In *Proc. of EMNLP*.
- Derrac, J., and Schockaert, S. 2015. Inducing semantic relations from conceptual spaces: A data-driven approach to plausible reasoning. *Artificial Intelligence* 228:66–94.
- Distel, F. 2011. *Learning description logic knowledge bases from data using methods from formal concept analysis*. Ph.D. Dissertation, Dresden University of Technology.
- Dubois, D.; Prade, H.; Esteva, F.; Garcia, P.; and Godo, L. 1997. A logical approach to interpolation based on similarity relations. *International Journal of Approximate Reasoning* 17(1):1–36.
- Ecke, A., and Turhan, A. 2012. Role-depth bounded least common subsumers for EL+ and ELI. In *Proc. of Description Logics*.
- Esteva, F.; Garcia, P.; Godo, L.; and Rodríguez, R. 1997. A modal account of similarity-based reasoning. *International Journal of Approximate Reasoning* 16(3-4):235–260.
- Faruqui, M.; Dodge, J.; Jauhar, S. K.; Dyer, C.; Hovy, E. H.; and Smith, N. A. 2015. Retrofitting word vectors to semantic lexicons. In *Proceedings of NAACL-HLT*, 1606–1615.
- Gärdenfors, P. 2000. *Conceptual spaces: The geometry of thought*. MIT press.
- Gärdenfors, P. 2014. *The geometry of meaning: Semantics based on conceptual spaces*. MIT press.
- Goldsmith, J.; Lang, J.; Truszczyński, M.; and Wilson, N. 2008. The computational complexity of dominance and consistency in CP-nets. *Journal of Artificial Intelligence Research* 33:403–432.
- Goodman, N. 1955. *Fact, fiction, and forecast*. Harvard University Press.
- Guo, S.; Wang, Q.; Wang, L.; Wang, B.; and Guo, L. 2016. Jointly embedding knowledge graphs and logical rules. In *Proc. of EMNLP*, 192–202.
- Gutiérrez-Basulto, V., and Schockaert, S. 2018. From knowledge graph embedding to ontology embedding? an analysis of the compatibility between vector space representations and rules. In *Proc. of KR*.
- Ibáñez García, Y.; Gutiérrez-Basulto, V.; and Schockaert, S. 2020. Plausible reasoning about \mathcal{EL} -ontologies using concept interpolation. <https://arxiv.org/abs/2006.14437>.
- Jameel, S., and Schockaert, S. 2016. Entity embeddings with conceptual subspaces as a basis for plausible reasoning. In *Proceedings of the 22nd European Conference on Artificial Intelligence*, 1353–1361.
- Jameel, S.; Bouraoui, Z.; and Schockaert, S. 2017. Member: Max-margin based embeddings for entity retrieval. In *Proceedings of the 40th International ACM SIGIR Conference on Research and Development in Information Retrieval*, 783–792.
- Jiang, Y.; Tang, Y.; Wang, J.; and Tang, S. 2009. Reasoning within intuitionistic fuzzy rough description logics. *Inf. Sci.* 179(14):2362–2378.
- Jung, J. C.; Lutz, C.; and Wolter, F. 2020. Least general generalizations in description logic: Verification and existence. In *Proc. of AAAI-20*.
- Klein, M. C. A.; Mika, P.; and Schlobach, S. 2007. Rough description logics for modeling uncertainty in instance unification. In *Proc. of URSW*.
- Kulmanov, M.; Liu-Wei, W.; Yan, Y.; and Hoehndorf, R. 2019. EL embeddings: Geometric construction of models for the description logic EL++. In *Proc. of IJCAI*.
- Küstners, R., and Borgida, A. 2001. What’s in an attribute? consequences for the least common subsumer. *J. Artif. Intell. Res.* 14:167–203.
- Li, N.; Bouraoui, Z.; and Schockaert, S. 2019. Ontology completion using graph convolutional networks. In *Proceedings of the 18th International Semantic Web Conference*, 435–452.
- Lieto, A., and Pozzato, G. L. 2018. A description logic of typicality for conceptual combination. In *Proceedings of the 24th International Symposium on the Foundations of Intelligent Systems*, 189–199.

- Lisi, F. A., and Straccia, U. 2013. A system for learning GCI axioms in fuzzy description logics. In *Description Logics*, volume 1014 of *CEUR Workshop Proceedings*, 760–778. CEUR-WS.org.
- Lutz, C., and Milicic, M. 2007. A tableau algorithm for description logics with concrete domains and general TBoxes. *J. Autom. Reasoning* 38(1-3):227–259.
- Lutz, C., and Wolter, F. 2011. Foundations for uniform interpolation and forgetting in expressive description logics. In *Proceedings of the 22nd International Joint Conference on Artificial Intelligence (IJCAI-11)*.
- Lutz, C. 2002. *The complexity of description logics with concrete domains*. Ph.D. Dissertation, RWTH Aachen University, Germany.
- Miclet, L.; Bayouh, S.; and Delhay, A. 2008. Analogical dissimilarity: definition, algorithms and two experiments in machine learning. *Journal of Artificial Intelligence Research* 32:793–824.
- Mikolov, T.; Chen, K.; Corrado, G.; and Dean, J. 2013. Efficient estimation of word representations in vector space. In *International Conference on Learning Representations*.
- Neelakantan, A., and Chang, M. 2015. Inferring missing entity type instances for knowledge base completion: New dataset and methods. In *Proc. of NAACL*, 515–525.
- Osherson, D. N.; Smith, E. E.; Wilkie, O.; Lopez, A.; and Shafir, E. 1990. Category-based induction. *Psychological review* 97(2):185–200.
- Özçep, Ö. L., and Möller, R. 2013. Spatial semantics for concepts. In *In Proc. of DL*.
- Peñaloza, R., and Zou, T. 2013. Roughening the envelope. In *Proc. of FroCos*.
- Porello, D.; Kutz, O.; Righetti, G.; Troquard, N.; Galliani, P.; and Masolo, C. 2019. A toothful of concepts: Towards a theory of weighted concept combination. In *Proc. of DL*.
- Prade, H., and Richard, G. 2014. From analogical proportion to logical proportions: A survey. In *Computational Approaches to Analogical Reasoning: Current Trends*. Springer. 217–244.
- Quine, W. 1953. *From a Logical Point of View*. Harvard University Press.
- Rizzo, G.; d’Amato, C.; Fanizzi, N.; and Esposito, F. 2013. Assertion prediction with ontologies through evidence combination. In *International Workshop on Uncertainty Reasoning for the Semantic Web*, 282–299.
- Rudolph, S. 2006. *Relational exploration: combining description logics and formal concept analysis for knowledge specification*. Ph.D. Dissertation, Dresden University of Technology, Germany.
- Schlobach, S.; Klein, M. C. A.; and Peelen, L. 2007. Description logics with approximate definitions - precise modeling of vague concepts. In *Proc. of IJCAI*.
- Schockaert, S., and Prade, H. 2013. Interpolative and extrapolative reasoning in propositional theories using qualitative knowledge about conceptual spaces. *Artificial Intelligence* 202:86–131.
- Sertkaya, B. 2010. A survey on how description logic ontologies benefit from FCA. In *Proc. of CLA*.
- Sheremet, M.; Tishkovsky, D.; Wolter, F.; and Zakharyashev, M. 2005. Comparative similarity, tree automata, and diophantine equations. In *Proc. of LPAR-2005*.
- Sheremet, M.; Wolter, F.; and Zakharyashev, M. 2010. A modal logic framework for reasoning about comparative distances and topology. *Ann. Pure Appl. Logic* 161(4):534–559.
- Straccia, U. 2001. Reasoning within fuzzy description logics. *Journal of Artificial Intelligence Research* 14:137–166.
- Tversky, A. 1977. Features of similarity. *Psychological review* 84(4):327.
- van Harmelen, F., and ten Teije, A. 2019. A boxology of design patterns for hybrid learning and reasoning systems. *J. Web Eng.* 18(1-3):97–124.
- Wille, R. 1982. Restructuring lattice theory: An approach based on hierarchies of concepts. In *Ordered Sets*. Springer Netherlands. 445–470.
- Xie, R.; Liu, Z.; Jia, J.; Luan, H.; and Sun, M. 2016. Representation learning of knowledge graphs with entity descriptions. In *Proceedings of the Thirtieth AAAI Conference on Artificial Intelligence*, 2659–2665.
- Xu, C.; Bai, Y.; Bian, J.; Gao, B.; Wang, G.; Liu, X.; and Liu, T.-Y. 2014. RC-net: A general framework for incorporating knowledge into word representations. In *Proc. of CIKM*.
- Zarrieß, B., and Turhan, A. 2013. Most specific generalizations w.r.t. general EL-TBoxes. In *Proc. of IJCAI-12*.